ABSTRACT

A file format for a serverless distributed file system is composed of two parts: a primary data stream and a metadata stream. The data stream contains a file that is divided into multiple blocks. Each block is encrypted using a hash of the block as the encryption key. The metadata stream contains a header, a structure for indexing the encrypted blocks in the primary data stream, and some user information. The indexing structure defines leaf nodes for each of the blocks. Each leaf node consists of an access value used for decryption of the associated block and a verification value used to verify the encrypted block independently of other blocks. In one implementation, the access value is formed by hashing the file block and encrypting the resultant hash value using a randomly generated key. The key is then encrypted using the user's key as the encryption key. The verification value is formed by hashing the associated encrypted block using a one-way hash function. The file format supports verification of individual file blocks without knowledge of the randomly generated key or any user keys. To verify a block of the file, the file system traverses the tree to the appropriate leaf node associated with a target block to be verified. The file system hashes the target block and if the hash matches the access value contained in the leaf node, the block is authentic.